BEACON SPACECRAFT OPERATIONS: LESSONS IN AUTOMATION

ROBERT SHERWOOD, ALAN SCHLUTSMEYER, MILES SUE, JOHN SZIJJARTO, E. JAY WYATT
Jet Propulsion Laboratory, California Institute of Technology

ABSTRACT

A new approach to mission operations has been flight validated on NASA's Deep Space One (DS1) mission that launched in October 1998. The beacon monitor operations technology is aimed at decreasing the total volume of downlinked engineering telemetry by reducing the frequency of downlink and the volume of data received per pass. Cost savings are achieved by reducing the amount of routine telemetry processing and analysis performed by ground staff, and by the reduced utilization of the communications ground stations. With beacon monitoring, the spacecraft assesses its own health and transmits one of four sub-carrier frequency tones to inform the ground of how urgent it is to track the spacecraft for telemetry. If all conditions are nominal, the tone provides periodic assurance to ground personnel that the mission is proceeding as planned without having to receive and analyze downlinked telemetry. If there is a problem, the tone will indicate that tracking is required and the resulting telemetry will contain a concise summary of what has occurred since the last telemetry pass.

KEYWORDS: Spacecraft Operations, Automation, Beacon

INTRODUCTION

The budget environment that has evolved since the advent of NASA's Faster, Better, Cheaper initiative has caused mission risk policies and mission designs to change in ways that have been conducive to the inception of new operations concepts and supporting technologies. Such was the case when the beacon monitor concept was conceived to enable a mission to Pluto to be achieved within the budget constraints passed down from NASA. The technology was accepted into the New Millennium Program and baselined for flight validation on the Deep Space One Mission. As the technology was being developed for DS1, the NASA community has expressed a growing interest and acceptance of adaptive operations and onboard autonomy.

In traditional mission operations, the spacecraft receives commands from the ground and in turn transmits telemetry in the form of science and/or engineering data. With beacon monitoring, the spacecraft sends a command to the ground that instructs the ground personnel how urgent it is to track the spacecraft for telemetry. There are only four such commands. Our approach is one where telemetry is only transmitted when it is necessary for ground personnel to assist the spacecraft. When telemetry tracking is necessary the intelligent data summaries contain the most relevant information to provide full insight into spacecraft activities since the last contact. The key challenge has been to develop an architecture that enables the spacecraft to adaptively create summary information to make best use of the available bandwidth as the mission progresses such that all pertinent data is received in one telemetry pass lasting four to eight hours.

The primary components of the technology are a tone messaging system, AI-based software for onboard engineering data summarization, a ground visualization system for telemetry summaries, and a ground response system. Beacon tone operations can be used to lower the cost of operating space missions while simultaneously decreasing their risk. The concept involves a paradigm shift from routine telemetry downlink and ground analysis to on-board health determination and autonomous data summarization. Beacon operations will enable more of the smaller, more frequent missions that NASA is planning for the early part of the next millennium. This paper includes a description of the Beacon monitor concept; the trade-offs associated with adapting that concept as a technology experiment, and our lessons learned during the DS1 mission. Applicability to future missions is also included.

DS1 BEACON EXPERIMENT SUBSYSTEMS

It was required that two subsystems be designed and developed to implement the desired functionality for the DS1 Beacon experiment. These are, in fact, standalone innovations. The two subsystems are the tone system and the data summarization system. The on-board software for these two systems contains three subroutines. Each of these subroutines will be described. The hardware for the ground tone detection is not discussed in this paper but is described in Reference 3. Although they are being presented here primarily in support of cruise phase operations on a deep-space mission, there has also been interest in applying these technology components to other domains. Other potential applications include using in-situ beacons at Mars, adapting tone messaging and summarization to Earth-orbiting spacecraft, using beacons for science event detection and notification, and in utilizing the tone system to reduce mission risk due to spacecraft operability constraints.

Tone System

There are four tone signals and each uniquely represents one of the four urgency-based beacon messages. The DS1 tone definitions are nominal, interesting, important, and urgent. These tones are generated as the spacecraft software reacts to real-time events. The definitions of these tones as used for DS1 are contained in Table 1.

During the DS1 tone experiment, the Beacon tone was sent at prescheduled times. The tone is set when anomalies are detected by the data summarization component of Beacon. We are currently testing routine operational use of the beacon monitor system during the DS1 extended mission, which began in September 1999.

Data Summarization System

Whenever the beacon tone indicates a need for tracking, the on-board summarization system provides concise summaries of all pertinent spacecraft data. The summarization system performs three functions: data collection and processing, mission activity determination, and episode identification.

Table 1. Beacon Tone Definitions

Tone	Definition
Nominal	Spacecraft is nominal; all functions are performing as expected. No need to downlink engineering telemetry.
Interesting	An interesting and non-urgent event has occurred on the spacecraft. Establish communication with the ground when convenient. Examples: device reset to clear error caused by single event upset due to cosmic particle, other transient events.
Important	Communication with the ground needs to be achieved within a certain time or the spacecraft state could deteriorate and/or critical data could be lost. Examples: memory near full, non-critical hardware failure.
Urgent	Spacecraft emergency. A critical component of the spacecraft has failed. The spacecraft cannot autonomously recover and ground intervention is required immediately. Examples: Propulsion or power system electronics failure
No Tone	Beacon mode is not operating, spacecraft telecom is not Earth-pointed or spacecraft anomaly prohibited tone from being sent.

The data collection subroutine receives data from the DS1 sensors and applies standard summary functions to this data. The functions are minimum, maximum, mean, first derivative, and second derivative. Performance summaries are generated at regular intervals and stored in memory until the next telemetry ground contact. By summaries every 15-30 minutes instead of raw engineering measurements every 1-5 seconds as is common for raw spacecraft data, significant reductions in communications bandwidth can be obtained.

The coarse summary data is useful for long-term trend analysis, but when anomalies occur it is necessary to have higher fidelity data. The episode subroutine has two inputs:

- Summary and raw sensor data received internally from the data collection subroutine. (The types of engineering data used on DS1 are listed in Figure 1.)
- A table of high and low limits based on the current mission mode from the Mission Activity Subroutine. (The mission modes are listed in Figure 1.)

The Mission Activity Subroutine is a classifier that takes the current summaries for each sensor and determines the current spacecraft mission mode. This mission mode will determine which limit table should be applied in the Episode Subroutine. When the raw sensor data or summarized data violate the high or low limits, the subroutine spawns an "episode" and outputs past relevant data from the data collection subroutine. The past data collected are one-minute summaries that go back in time several minutes before the episode started. The episode subroutine outputs this data to the telemetry subsystem for downlink. In addition, the episode subroutine will change the state of the beacon tone based on the severity of the anomaly. A flow chart of the data summarization software is contained in Figure 1.

The software also has the capability to use AI-based envelope functions instead of traditional alarm limits. This system, called ELMER (Envelope Learning and Monitoring using Error Relaxation), provides a new form of event detection and will be evaluated in addition to using the project-specified traditional alarm limits. Envelope functions are essentially adaptive alarm limits learned by training a neural network with nominal engineering data. The neural net can be on-board or on the ground. For DS1, envelope functions are trained on the ground and then uploaded to the spacecraft.

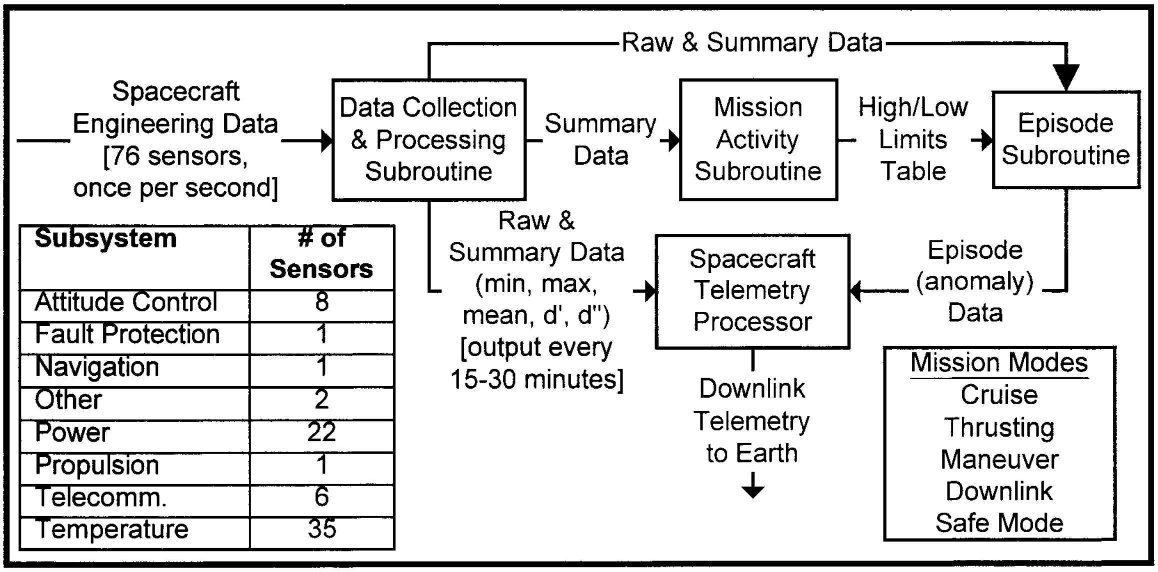

Subsystem	# of Sensors
Attitude Control	8
Fault Protection	1
Navigation	1
Other	2
Power	22
Propulsion	1
Telecomm.	6
Temperature	35

Figure 1. Data Summarization Software Flow Chart

RESULTS & LESSONS LEARNED

A series of experiments were run to test the end-to-end tone delivery system. These experiments were designed to incrementally test additional capability for the Beacon tone system. Prior to launch, the ability of the Small Deep Space Transponder (SDST) to generate Beacon tones was tested by the telecom engineers. A similar test was performed on the spacecraft several times after launch. Additional tests were run with weak signals, and using a different frequency transmitter.

The data summarization component of Beacon was activated in February 1999. The Beacon Team determined the limits applied to the engineering data for testing the summarization capability. The limits were set just outside of the minimum and maximum value seen for the data since launch. Shortly after the first turn-on several of the data channels went into episode (out-of-limits) condition. Upon further inspection, it was determined that many limits were based on engineering units (EU), but much of the data was being stored using data numbers (DN) in the on-board engineering and housekeeping telemetry system (EH&A). To correct this problem, the initialization file was updated with DN based limits and uplinked.

Additional tests with the data summarization software have been successful. There have been a few false alarms due to DS1 flight software bugs, and because of transient effects. Beacon data summarization has been an evolving process requiring several limit refinements from the spacecraft team. This should be expected in the development of any data summarization system. This process is very similar for any new mission launch. For the first several months, ground alarms are updated as the flight team learns about how the spacecraft really operates. The ground testing activities give a good first cut at setting alarm levels, but the spacecraft never operates exactly as it did in test. Implementing context sensitive limits is a similar process. We are no longer setting the engineering data limits based on the worse case. Now we can look at the worse case based on the spacecraft activities. This should ensure more accurate anomaly discovery.

The DS1 project has been using Beacon to periodically check spacecraft health during the extended mission. As a result, fewer telemetry downlink passes have been scheduled. This has resulted in a cost savings and additional resources available for other projects. In addition, data summarization has been used during periods of low bandwidth communications to provide additional visibility into spacecraft performance.

Design Lessons

The utilization of the ion propulsion system (also called solar-electric propulsion) on DS1 offers an additional advantage in using Beacon monitoring. The IPS provides continuous thrust for much of the cruise phase. The operational margin for IPS thrusting represents the duration for which IPS could be off and still allow the spacecraft to reach the target asteroid. Due to the low thrust associated with IPS and because actual thrusting did not start until several weeks after launch, the operational margin is only a few weeks. Telemetry downlink passes became less frequent as the DS1 mission progressed. Currently, there is only one telemetry pass per week. If the spacecraft experiences a problem that requires the standby mode, the IPS engine will be shut down. It could be up to one week before the flight team knows the spacecraft is in standby mode. Using the Beacon tone system during the periods between scheduled telemetry downlinks can be a cost-effective way to decrease mission risk because it reduces the likelihood of losing thrusting time and not making the intended target. Other future IPS missions have taken note of this fact and requested Beacon tone services to lower their mission risk.

Implementation Lessons

Project management decided to redesign the DS1 flight software about 18 months before launch. This decision greatly compacted an already full schedule to complete the software. As a result, the testing of all non-essential software functions was delayed until after launch. The Beacon experiment was considered a non-essential piece of software and therefore was only tested pre-launch for non-interference with the other flight software. In post launch testing, a few problems were discovered that prevented us from starting the Beacon software until a new version could be uploaded. These problems related to differences between the flight hardware-based testbed and a simulated hardware testbed. This is the age-old lesson learned of performing system testing on the software prior to use. But even beyond that, it is important to run tests on the actual hardware-based testbed. Unfortunately, the DS1 schedule forced us to wait until post launch.

Before the software redesign, the Beacon software was tightly integrated with the DS1 fault protection software. To reduce risk, project management decided after the redesign to de-couple the two pieces of software. Previously, the fault protection monitors triggered the Beacon tones. After the redesign, the mapping of faults to tones was performed using two different methods. All spacecraft standby modes are now mapped to the urgent Beacon tone. The interesting and important Beacon tones are mapped using Beacon software determined limits. Decoupling the fault protection software from the Beacon software gives us maximum flexibility to determine what sensors to monitor. Unfortunately, our algorithms for determining faults are not nearly as sophisticated as the fault protection monitors. These monitors can look at many different values based on conditional logic before determining what fault has occurred. Future spacecraft designed to use Beacon operations should plan on completely integrating the Beacon tone software with the fault protection software.

Flight Test Lessons

During some of the DS1 tone experiments, the initial frequency uncertainty was much larger than expected. Temperature variations that occurred after optical navigation maneuvers caused this large uncertainty. A bias was manually introduced to keep the ground-received signal in the recorded band. Without the bias, the frequency might be outside the recorded band. In an automated detection mode, it is necessary to record at least 3 times the current bandwidth, unless a better way to predict the frequency can be found. One possibility is to make use of the Auxiliary Oscillator Frequency vs. Temperature calibration table to improve frequency prediction.

In one of the experiments, the actual tone switching times did not seem to agree exactly with the predicted switching times. This led to the DS1 team discovering an error of 18-19 seconds in the on-board spacecraft time to Earth receive time conversion.

We noticed problems with false episode alarms due to infrequent mission activities such as optical navigation manuevers, camera calibrations, etc. It is important to carefully define each of the mission modes and how they are related to the engineering data produced during these infrequent activities. In the DS1 case, we had defined the maneuver mission mode to only occur when the thrusters were firing. Since maneuvers also involved turning the spacecraft, it was important to include all events that turned the spacecraft in our maneuver mission activity criteria. Once mission modes are carefully defined, then episode limits for those modes can be developed.

Two tone passes were not successful due to the tracking station's (DSS-13) weather and equipment. In one experiment, the spacecraft started transmitting tones before it rose above the horizon of DSS13. In another case, a scheduled pass was cancelled due to spacecraft activities. While the overall tone experiments have been very successful, future experiment plans should allow for this kind of contingency.

CONCLUSIONS

Beacon operations is a valuable tool for reducing overall mission risk in an environment where decreased tracking is all but mandated by slim operations budgets. It can also be viewed as a technology for conducting low cost mission operations at acceptable risk. The key point here is that NASA policy towards mission risk and cost changed when the visions for smaller, faster, better, and cheaper missions were born. Beacon operations helps enable many more missions with existing tracking resources and is a practical method for minimizing mission risk while decreasing the frequency of telemetry tracking and staffing levels to save operational cost. The Beacon experiment on DS1 has proven the functionality of the technology. It has also shown that it can be effective in reducing downlink volume and frequency, summarizing spacecraft engineering telemetry, and reducing operations costs. Additional use of Beacon on DS1 should prove that Beacon operations and cost reductions are sustainable in the long-term. Future missions should be able to benefit from this proven technology.

ACKNOWLEDGEMENT

The research described in this paper was partly funded by the TMOD Technology Program and performed at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

REFERENCES

1. Sherwood, R., et. al. "Lessons Learned During Implementation and Early Operations of the DS1 Beacon Monitor Experiment." *Third International Symposium on Reducing the Cost of Ground Systems and Spacecraft Operations*. Tainan, Taiwan. (1999).
2. Wyatt, E. J., et. al. "An Overview of the Beacon Monitor Operations Technology." *International Symposium on Artificial Intelligence, Robotics, and Automation in Space*. Tokyo, Japan. (1997).
3. Wyatt, E. J., et. al. "Beacon Monitor Operations on the Deep Space One Mission." *Fifth International Symposium on Space Mission Operations and Ground Data Systems*. Tokyo, Japan. (1998).
4. DeCoste, D. "Automated Learning and Monitoring of Limit Functions." *International Symposium on Artificial Intelligence, Robotics, and Automation in Space*. Tokyo, Japan. (1997).
5. Sherwood, R., et. al. "Flight Software Implementation of the Beacon Monitor Experiment On the NASA New Millennium Deep Space 1 (DS1) Mission." *Second International Symposium on Reducing the Cost of Spacecraft Ground Systems and Operations*. Oxfordshire, UK. (1997).
6. Wyatt, E.J. and J.B. Carraway. "Beacon Monitoring Approach to Spacecraft Mission Operations." *First International Symposium on Reducing the Cost of Spacecraft Ground Systems and Operations*. Oxfordshire, UK. (1995).